Abstract Voronoi Diagrams from Closed Bisecting Curves *

Cecilia Bohler, Rolf Klein, and Chih-Hung Liu

University of Bonn, Institute of Computer Science I, D-53113 Bonn, Germany.

Abstract. We present the first algorithm for constructing abstract Voronoi diagrams from bisectors that are unbounded or closed Jordan curves. It runs in expected $O\big(s^2 n \log(\max\{s,n\}) \sum_{i=2}^{n} m_i/i\big)$ many steps and $O(\sum_{i=3}^{n} m_i)$ space, where m_i denotes the average number of faces per Voronoi region in any diagram of a subset of i sites.
Keywords: Abstract Voronoi diagrams, closed bisecting curves, computational geometry, distance problems, Voronoi diagrams.

1 Introduction

Abstract Voronoi diagrams are a unifying concept that provides structural results and efficient algorithms for all concrete Voronoi diagrams whose bisector systems satisfy some simple combinatorial properties. For example, if all bisectors are unbounded Jordan curves, if all Voronoi regions of three sites are pathwise connected and if their closures cover the plane, then it is known that the Voronoi diagram of n sites is a planar graph of linear complexity which can be constructed in expected $O(n \log n)$ many steps [4, 17, 15]. Quite a few new types of Voronoi diagrams were found to be under this roof, freeing authors from developing tailor-made algorithms [15].

The concept of abstract Voronoi diagrams has been extended to higher orders in [5, 8]. Linear algorithms for special bisector systems were presented in [16, 7]. Also, the above conditions on the bisecting curves were relaxed in [6], by allowing Voronoi regions to be disconnected. But so far, only unbounded bisecting curves were studied, leaving many interesting Voronoi diagrams uncovered.

In this paper we consider, for the first time, abstract Voronoi diagrams built from arbitrary Jordan curves, unbounded or closed. We require the following properties. Bisecting curves intersect vertical lines only constantly often. Any two bisecting curves intersect finitely often, and at most $\leq s$ times if both bisectors are related to the same site. The closures of the Voronoi regions of any three sites cover the plane. Under these assumptions we can construct the abstract Voronoi diagram in expected

$$O\Big(s^2 n \log(\max\{s,n\}) \sum_{i=2}^{n} \frac{m_i}{i}\Big) \tag{1}$$

* This work was partially supported by the European Science Foundation (ESF) in the EUROCORES collaborative research project EuroGIGA/VORONOI.

many steps and in $O(\sum_{i=3}^{n} m_i)$ space, where m_i denotes the average number of faces per Voronoi region in any diagram of a subset of i sites. Let us look at some applications.

If the bisecting curves are pseudo-circles, Voronoi regions are of linear complexity, as we will see in Section 5. Thus, $m_i \leq i$, and (1) yields the first $O(n^2 \log n)$ randomized algorithm for this class of Voronoi diagrams, which is nearly optimal since their size can be quadratic.

An important special case are points with multiplicative weights; here bisectors are lines or Euclidean circles. The deterministic algorithm by Aurenhammer and Edelsbrunner [3] runs in time $\Theta(n^2)$ which is optimal in the worst-case where the diagram is of quadratic size. In this case, our algorithm takes a $\log n$ factor longer. But if, e.g., points are fixed, and weights drawn independently from some distribution (or just permuted at random) then the multiplicatively weighted Voronoi diagram has expected size only $O(n \log^2 n)$, as was recently shown by Har-Peled and Raichel [14], and it can be constructed in this time. Our algorithm is sensitive to the smaller output size: Since $m_i \in O(\log^2 i)$ holds in this case, it runs in expected time $O(n \log^4 n)$.

For the farthest Voronoi diagram of n simple polygons of constant complexity, formula (1) yields a randomized $O(n \log^2 n)$ algorithm. Indeed, s is constant, and the Voronoi diagram is of linear size, as was shown by Cheong et al. [10]; thus, the average number of faces per Voronoi region is constant. In comparison, the algorithm in [10] runs on polygons of total complexity n in deterministic time $O(n \log^3 n)$.

Our algorithm for abstract Voronoi diagrams is fairly general. It uses randomized incremental construction [12, 11] based on a new history graph and search procedure that is quite different from [17, 15] and [6]. Further, it uses an application of Chazelles linear triangulation algorithm [9] to compute the trapezoidal decomposition of a given face. At a later stage this may be replaced by a simpler randomized method, e. g. by Amato et. al. [2], which though would require a new randomized analysis.

We will discuss the new history graph and search procedure in Section 3, and upper bound the expected number of conflicts, and the effort invested in finding them in Section 4. Because of space limitations some proofs and comments are put in the appendix.

2 Preliminaries

This section contains basic definitions and well-known facts about abstract Voronoi diagrams. We are given a set S of n virtual sites $p, q, r, \ldots$, and for any two sites $p \neq q$ in S a bisecting curve $J(p, q) = J(q, p)$. These curves are supposed to satisfy the following five requirements.

(A1) Under stereographic projection, each curve $J(p, q)$ is a closed Jordan curve on the sphere.

Curve $J(p,q)$ is unbounded iff its projection runs (in the limit) through the north pole of the sphere; otherwise, $J(p,q)$ is a closed Jordan curve in the plane. Each such curve, unbounded or closed, bisects the plane into two open domains. As part of the definition of $J(p,q)$, they are denoted by $D(p,q)$ and $D(q,p)$, and in drawings we write the site p on the $D(p,q)$-side of the bisector.

Intuitively, points in $D(p,q)$ are closer to virtual site p than to q. We define the *Voronoi region* of $p \in S$ by

$$\mathrm{VR}(p,S) := \bigcap_{q \in S \setminus \{p\}} D(p,q),$$

and the *abstract Voronoi diagram*, $V(S)$ of S, as the plane minus the union of all Voronoi regions.

(A2) For any three sites $p_i \in S$, where $i = 1,2,3$, the closures of the Voronoi regions $VR(p_i, \{p_1,p_2,p_3\})$ cover the plane.

A useful consequence of property (A2) is the following Transitivity Lemma [15].

Lemma 1. *For any p,q,r in S we have $D(p,q) \cap D(q,r) \subseteq D(p,r)$.*

This fact would be obvious if we could read "$z \in D(p,q)$" as "z is closer to p than to q". Here it has to be derived from axioms (A1) and (A2). Lemma 1 implies that for *every* subset $T \subseteq S$ the closures of the Voronoi regions $\mathrm{VR}(t,T)$ cover the plane [15].

To identify the graph structure of the point set $V(S)$, the following property is convenient.

(A3) The intersection of any two bisecting curves $J(p,q), J(r,s)$ has a finite number of connected components.

Property (A3) allows us to zoom in on any point v situated on one or more bisecting curves, and find a neighborhood U of v where the curves passing through v either stay disjoint or coincide, on either side of v. As a consequence one obtains a "piece-of-pie" Lemma stating that each point v of $V(S)$ is either an interior point of a Voronoi edge, if v lies on the common boundary of exactly two Voronoi regions of different sites, or neighborhood U contains pieces of Voronoi regions of three different sites or more, making v a Voronoi vertex. An example of an abstract Voronoi diagram based on A1–A3 is shown in Figure 1. The main differences to standard abstract Voronoi diagrams [15] are as follows. There can be islands inside a Voronoi region. The boundary of an island need not contain a Voronoi vertex and it equals a closed bisector $J(p,q)$. Voronoi regions can be disconnected. In the neighborhood of a Voronoi vertex, the same face of a Voronoi region can be represented several times. In order to bound the complexity of the Voronoi diagram we require the following.

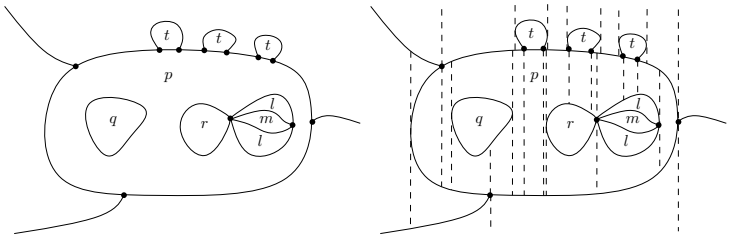

Fig. 1: An abstract Voronoi diagram based on axioms A1, A2, and A3 and its trapezoidal decomposition.

(A4) The intersection of any two bisecting curves $J(p,q), J(p,s)$ related to the same site p has at most s connected components.

Property A4 implies that, in a Voronoi diagram of three sites, the number of vertices and faces of a Voronoi region is in $O(s)$. For the overall complexity of the Voronoi diagram, one obtains from (A4) the following.

Lemma 2. *A single Voronoi region is of complexity $O(s\,n^2)$, and this bound can be attained. The whole Voronoi diagram $V(S)$ is of size $O(s\,n^3)$ [6].*

The proof is by vertex counting. We observe that islands can be nested, but their total number is in $O(n)$. Namely, if an island is situated in a face of the Voronoi region of p, and contains points of the region of some site $q \neq p$, then the whole Voronoi region of q is confined to this island by the bisector $J(q,p)$; see Figure 2a.

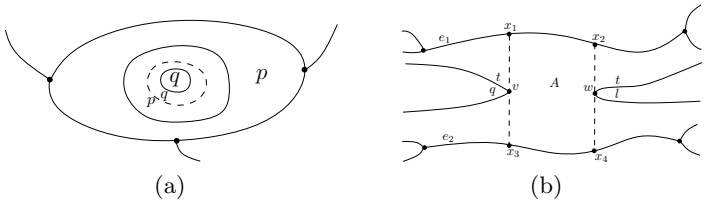

Fig. 2: (a) The Voronoi region of q is confined to the island by the dashed bisector $J(q,p)$. (b) The description of trapezoid A.

Since we are going to subdivide Voronoi regions into trapezoids [18] we require one last property which is quite common in arrangement theory [19].

(A5) Each bisecting curve has at most a constant number of points of vertical tangency. All such points have pairwise different x-coordinates.

Property A5 implies that each bisecting curve has at most a constant number of points of vertical tangency. The trapezoidal decomposition $V^*(R)$, for a subset R of S, results from shooting vertical rays from all vertices and points of vertical tangency in both directions, until they hit another point of $V(R)$; see Figure 1. Note that $V^*(R)$ is still of the same size as $V(R)$.

Voronoi edges and trapezoids of $V^*(R)$ are referred to by unique *descriptions*. A general edge description is shown in the leftmost drawing of Figure 3. Associated with vertex v is a 5-tuple containing the names of the two sites p, q separated by the edge, the names of the adjacent sites, r_q and r_p, and the coordinates of v. The sites are listed in clockwise order around v. A similar 5-tuple is associated with w.

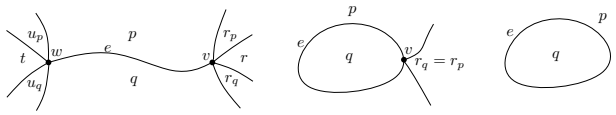

Fig. 3: Edges e with descriptions $D_R(e) = \{(r_q, q, p, r_p, \text{coor}(v)), (u_p, p, q, u_q, \text{coor(w)})\}$, $D_R(e) = \{(r_q, p, q, r_p), \text{coor}(v)\}$ and $D_R(e) = \{(p, q)\}$.

Figure 2b depicts a general trapezoid. Its description consists of the descriptions of the upper and lower Voronoi edges e_1 and e_2, the sites q and l whose bisectors $J(q,t)$ and $J(l,t)$ determine the tangency points v and w, and the coordinates of x_1 to x_4.

As a *basic operation* we assume that the intersections of two related bisectors can be computed in time $O(s)$, see (A4). The intersection between a Voronoi edge of $V(R)$ and the region of a new site $\text{VR}(t, R \cup \{t\})$ for $t \in S \setminus R$ can also be computed in time $O(s)$, because this can be decided in the diagram of 7 sites, the site t and the sites contained in the description of e, and e can be intersected in at most $O(s)$ components. The intersection of a vertical line and a bisector can be computed in constant time, see (A5), and hence, the intersection of a trapezoid and a new region can be computed in time $O(s)$. Further, we assume that for a bisector $J(p,q)$ and a point x in the plane, we can in constant time decide if x is on $J(p,q)$ or in $D(p,q)$ or $D(q,p)$. Also, for two points x and y on $J(p,q)$, given a direction of $J(p,q)$, we can decide in constant time if x or y comes first.

3 Searching for Intersections

For the algorithm we restrict our attention to the finite part of the Voronoi diagram and assume that a large closed curve Γ is given, encircling the diagram such that it contains all closed bisectors, it intersects each unbounded bisecting curve exactly twice, and no pair of bisecting curves intersect on or outside of Γ.

Our algorithm is a randomized incremental construction, where the insertion order $\{r_1, \ldots, r_n\} = S$ of the sites is given randomized, and we insert the Voronoi

regions successively to the current Voronoi diagram. Let $R = \{r_1, \ldots, r_j\}$, $V^*(R)$ the trapezoidal decomposition of the current diagram $V(R)$, and $t \in S \setminus R$ the next site to be inserted. In order to incrementally construct the augmented Voronoi diagram $V^*(R \cup \{t\})$ from $V^*(R)$ we need to determine which parts of $V^*(R)$—edge or trapezoid—are intersected by the Voronoi region $\mathrm{VR}(t, R \cup \{t\})$ of the new site t (if this region is a small island it could well be contained in the interior of a single trapezoid of $V^*(R)$, without intersecting any edges).

We define edges, trapezoids, and Voronoi regions to be open, i.e., an edge does not contain its endpoints, and trapezoids and regions do not contain their boundaries. Thus, all intersections are "proper", in that they contain a subset of full dimension. If q were the last site inserted to the Voronoi diagram shown in Figure 1, its region would not intersect the previous Voronoi diagram, but one of its trapezoids.

Definition 1. *We say that a trapezoid A or an edge e of $V^*(R)$ is*

(i) intersected *by t, if $VR(t, R \cup \{t\})$ intersects A or e respectively;*
(ii) in conflict *with t, if A or e with their descriptions $D_R(A)$ and $D_R(e)$ do not appear in $V^*(R \cup \{t\})$.*

Observe that if a trapezoid or an edge is intersected by t, then it is also in conflict with t but not necessarily vice versa, see Figure 4a. Recall also the very important fact that if $R' \subseteq R$, then $\mathrm{VR}(t, R \cup \{t\}) \subseteq \mathrm{VR}(t, R' \cup \{t\})$, which follows quite directly from the definition of Voronoi regions.

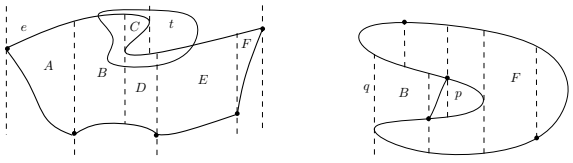

(a) (b)

Fig. 4: (a) All trapezoids A to F are in conflict with t, because t intersects edge e, but only the trapezoids B to E are intersected by t. (b) Trapezoid B borders on two different edges of the face F.

Finding all proper intersections of $\mathrm{VR}(t, R \cup \{t\})$ with $V^*(R)$ is facilitated by a *history graph* $\mathcal{H}(R)$, a DAG with a single source, a vertex for each face of the inserted regions, and a vertex for each trapezoid ever constructed during the incremental process. More precisely, let $R = R_j = \{r_1, \ldots, r_j\}$ be the insertion order, then the vertex set of $\mathcal{H}(R)$ equals

$$\{\text{source}\} \cup \{D_{R_i}(A) \mid A \text{ is a trapezoid of } V^*(R_i) \text{ for } i \in \{2, \ldots, j\}\}$$
$$\cup\ \{F \mid F \text{ is a face of } \mathrm{VR}(r_i, R_i) \text{ for } i \in \{3, \ldots, j\}\}.$$

Note that a trapezoid that exists in a sequence of Voronoi diagrams $V^*(R_j)$ is represented in $\mathcal{H}(R)$ by a single vertex. Faces are represented by a vertex only when contained in a new region, thus, in contradistinction to trapezoids, not all faces ever constructed are represented by a vertex in $\mathcal{H}(R)$. Let a vertex of $\mathcal{H}(R)$ be called *face-vertex* if it refers to a face and *trapezoid-vertex* if it refers to a trapezoid. The edges of $\mathcal{H}(R)$ are defined incrementally, maintaining the invariant that the leaves of $\mathcal{H}(R)$ correspond to the trapezoids of $V^*(R)$. Also each trapezoid and face of the current diagram $V^*(R)$ is linked to its corresponding trapezoid- and face-vertex of $\mathcal{H}(R)$ and vice versa.

The top of the history graph is an artificial source and its children are the trapezoids of the diagram of the first two sites r_1 and r_2 of the insertion order. Now let $t \in S \setminus R$ be the next site to be inserted. The history graph $\mathcal{H}(R \cup \{t\})$ is obtained by updating $\mathcal{H}(R)$ as follows.

If F is a face of $\mathrm{VR}(t, R \cup \{t\})$, then F is made a successor of all trapezoids A' of $V^*(R)$ (these trapezoids are leaves of $\mathcal{H}(R)$) intersected by F, compare face F in Figure 5.

If A is a trapezoid of $V^*(R \cup \{t\})$, which has not yet been part of $V^*(R)$, then there are two cases:

1. Trapezoid A is contained in an r-region for an old site $r \in R$. Then A is made a successor of all trapezoids A' of $V^*(R)$ intersected by A. Compare trapezoids $B_1, D_1, E_1, H_1, H_2, H_3, I_1, J_1$ in Figure 5.
2. Trapezoid A is contained in a face F of $\mathrm{VR}(t, R \cup \{t\})$. Then A is made a successor of F, compare trapezoids K, L, M, N, O, P in Figure 5.

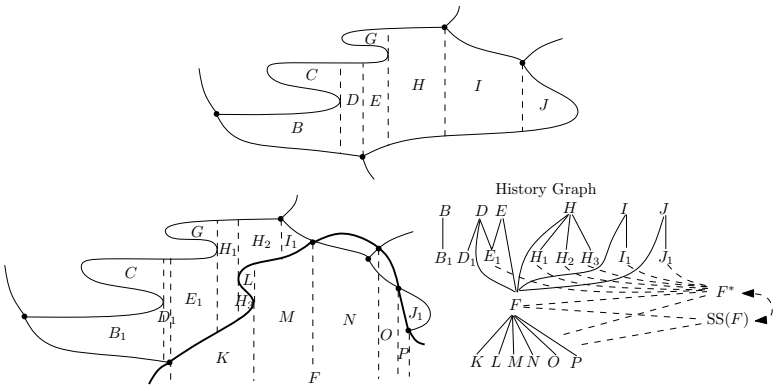

Fig. 5: Face $F \subseteq \mathrm{VR}(t, R \cup \{t\})$ is inserted in the trapezoidal decomposition of the upper diagram and the history graph is updated.

In addition to the history graph which contains only the descriptions of the trapezoids ever constructed and the "names" of the new faces, for each face-vertex F of $\mathcal{H}(R)$ we store its geometric structure. More specific, let F be

a face of $\mathrm{VR}(t, R \cup \{t\})$. We store its trapezoidal decomposition F^* together with a *search-structure* $\mathrm{SS}(F)$ which allows us to perform fast *point location* for a given query point x of the plane. We link the face-vertex F of $\mathcal{H}(R)$ to its geometric structures F^* and $\mathrm{SS}(F)$, and the trapezoids of F^* and $\mathrm{SS}(F)$ to their corresponding trapezoid-vertices, which are successors of F. Further, for each trapezoid A of $V^*(R \cup \{t\})$ bordering on an edge e of the boundary of F (this trapezoid must have been constructed while inserting F to $V^*(R)$), we link the trapezoid-vertex A to the edge e in F^*, and vice versa, compare trapezoids $E_1, H_1, H_2, H_3, I_1, J_1$ in Figure 5. Note that trapezoid A may border on up to two different edges of the boundary of F, compare Figure 4b. If this is the case, we link A to both of them.

Now let us discuss how to compute the geometric structures of F from scratch. Each edge on the boundary of F can be split up into a constant number of x-monotone arcs, by putting an additional vertex on each point of vertical tangency (A5). By doing this, in time $O(|F|)$, we obtain a simple polygon with monotone curved edges whose trapezoidal decomposition can be computed in time $O(|F|)$, see [9]. Observe that we consider F without the rest of the diagram and we compute the trapezoidal decomposition both in the inside and outside of F. The trapezoids in the outside are needed to obtain a monotone subdivision F^* of the whole plane. For this subdivision we can now built a *search-structure* $\mathrm{SS}(F)$, based on the structure for monotone subdivisions introduced by Edelsbrunner et. al. in [13], which again takes time and storage $O(|F|)$, and it allows us to perform *point location* for a given query point x of the plane in time $O(\log |F|)$.

Both in [9] and [13] the edges are straight lines, but the same ideas can be adapted to instances where edges are monotone curves, and the intersection between an edge and a vertical line can be computed in constant time.

Now we describe how to find all trapezoids of $V^*(R)$ intersected by the region of the new site t. We walk through the history graph along its trapezoid-vertices intersected by t as follows, starting with the successors of the source. If a trapezoid-vertex A of $V^*(R_i)$ is intersected by $\mathrm{VR}(t, R_i \cup \{t\})$, then we test recursively all its succeeding trapezoid-vertices for intersection with t; for the succeeding face-vertices we discuss later how to proceed. First, if A is linked to an edge e of the trapezoidal decomposition F^* of a face-vertex F, we test if e is intersected by t and if yes, we test all trapezoids bordering on e from the inside of F for intersection with t; these trapezoids are found in F^*. Afterwards we set a flag on e saying that e already has been tested for intersection with t, thus if e is reached from another trapezoid, we do not test it and its adjacent trapezoid for intersection with t again. For the running time analysis later on we note already now that if edge e is intersected by t, then all trapezoids bordering on e (there may be many!) are in conflict with t.

For each trapezoid bordering on e from the inside of F, having been successfully tested for intersection with t, we recursively test also its at most 4 neighbors in F for intersection with t. Doing this we trace the region of t through F. Of course the region of t may intersect F and other faces of the same region in several disconnected components, but we will later see that we will find a starting point

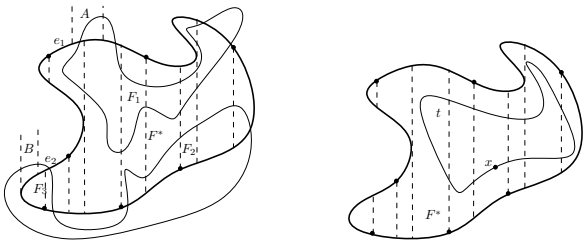

Fig. 6: The region of t (faces with normal boundary) is traced through the trapezoidal decomposition F^* (face with fat boundary), in the left figure from two starting points – one from A trough e_1, and one from B through e_2, and in the right figure from a point x found using $\mathrm{SS}(F)$.

for tracing each component. Again we set a flag with respect to t on each tested trapezoid to prevent it from having to be tested again.

In the left drawing of Figure 6, the region of t intersects F^* in 3 components, F_1, F_2, and F_3. If we enter F^* from A through e_1, then all trapezoids of F^* intersected by F_1 and F_2 are found. To detect the trapezoids intersected by F_3, we need another starting point, e. g. from B through e_2.

We may now have detected additional trapezoids which are intersected by t, they are linked to their corresponding trapezoid in the history graph and we recursively test also their successors for intersection with t.

For a face-vertex $F_{r_i} \subseteq \mathrm{VR}(r_i, R_i)$ reached in the history graph, we choose an arbitrary point x of the bisector $J(t, r_i)$ and use $\mathrm{SS}(F_{r_i})$ to determine if x is contained in F and if yes, in which trapezoid. If no trapezoid of F_{r_i} containing x is detected, i. e. x lies outside of F_{r_i}, then nothing has to be done, as will become clear later. Otherwise a trapezoid $A \subseteq F_{r_i}$ is detected with $x \in \overline{A}$. Then we know that A is intersected by t, and like before we can trace the region of t through F_{r_i} and mark the tested trapezoids with flags with respect to t (unless they already have a flag with respect to t, then we are done), compare the right drawing of Figure 6. Again we may have detected additional trapezoids intersected by t and we recursively test also their successors in the history graph for intersection.

Lemma 3. *By walking through the history graph $\mathcal{H}(R)$ as described above, we reach all trapezoids of $V^*(R)$ which are intersected by $VR(t, R \cup \{t\})$.*

Proof. (See first proof in Appendix) We prove that if a trapezoid A of $\mathcal{H}(R)$ is intersected by t, then either a predecessor of A in $\mathcal{H}(R)$ is intersected by t, or A is detected in the search structure $\mathrm{SS}(F)$ of a face F, or found while tracing the region of t through a face F. In all these cases, we will detect A in the history graph. □

Let E_t be the set of all trapezoids of $V^*(R)$ intersected by t. Once this set has been computed, we can update $V^*(R)$ like described in [6] and $\mathcal{H}(R)$ like described above. This takes time $O(s^2|E_t|)$, see [6].

4 Analysis

To organize the discussion of the running time of our algorithm we define the following variables. Like before let $R_j = \{r_1, \ldots, r_j\} \subseteq S$, the sites are inserted in this order, and $t = r_{j+1} \in S \setminus R_j$ is the next site to be inserted. Further

$$E_t := \{A \mid A \text{ is a trapezoid of } V^*(R) \text{ intersected by } t\}, \text{ and}$$

$$c := \sum_{i=3}^{j} \#\text{ trapezoids of } V^*(R_i) \text{ not in } V^*(R_{i-1}) \text{ in conflict with } t.$$

Lemma 4. *The set E_t can be computed in time $O(s^2 c \log(\max\{s, n\}))$.*

For the following randomized analysis we use the techniques by Clarkson, Mehlhorn, Seidel [11]. The proofs are similar to the ones in [6], for completion the missing ones can be found in the Appendix. To upper bound the number of conflicts during the insertion process in the algorithm, we define m_i as the average number of faces per region, over all diagrams of i sites from S.

Lemma 5. *For $R = \{r_1, \ldots, r_j\}$, the expectation of c is $O(\sum_{i=2}^{j} \frac{m_i}{i})$.*

Lemma 6. *The expected size of the history graph $\mathcal{H}(S)$ is $O(\sum_{i=3}^{n} m_i)$.*

Lemma 7. *The expected size of E_t is $O(m_j)$.*

Theorem 1. *Let $\{J(p,q) : p \neq q \in S\}$ be a bisecting curve system fulfilling axioms (A1) to (A5). Then $V(S)$ can be computed in expected time*

$$O\Big(s^2 n \log(\max\{s, n\}) \sum_{i=2}^{n} \frac{m_i}{i}\Big), \tag{2}$$

and expected space $O(\sum_{i=3}^{n} m_i)$.

Proof. Because of Lemma 4 and 5, the sets E_{r_i}, $3 \leq i \leq n$ can be computed in the claimed expected time, formula (2). Afterwards $\mathcal{H}(S)$ and $V^*(S)$ can be computed in time $O(s^2 \sum_{i=3}^{n} |E_{r_i}|)$, see the end of the previous section, and because the expectation of $|E_{r_i}|$ is $O(m_i)$, the expected value is in $O(s^2 \sum_{i=2}^{n} m_i) \in (2)$.

Lemma 6 shows the expected size of $\mathcal{H}(S)$, and the size of the trapezoidal decomposition and the search structures for the faces F is linear in the size of F, see [13], which gives us the expected space. □

5 Pseudo-circles: an Application

Let the set of bisectors $\mathcal{J} := \{J(p,q) : p \neq q \in S\}$ be a set of pseudocircles. Thus, each bisector is a simply closed curve and any two bisectors $J(p,q)$ and $J(r,t)$ are either equal, or their intersection is empty, exactly one nontransversal point, or exactly two transversal points. Such a curve system fulfills axioms (A1),

(A3), and (A4), in addition we assume that also (A2) and (A5) are fulfilled. If $|S| = 3$ then all Voronoi regions are connected, but in contradistinction to Voronoi diagrams with only unbounded bisectors, this does not imply that all regions in $V(S)$ are connected, see the Appendix for more details.

Nevertheless, we have the following complexity bound on the size of the Voronoi diagram. Agarwal et. al. [1] showed that the boundary of the union of n pseudodiscs, for $n \geq 3$, consist of at most $6n - 12$ elementary arcs. In our case we are interested in the complexity of the intersection of pseudodiscs and complements of pseudodiscs, thus the result by Agarwal et. al. is not directly applicable, but the technique used in the last part of our proof is similar.

Theorem 2. *Let S be a set of n sites together with a set of bisecting curves fulfilling our axioms. Then each Voronoi region in $V(S)$ is of size $O(n)$ and $V(S)$ is of size $O(n^2)$. These bounds are tight in the worst case.*

Proof. (Sketch) Let p be a site in S. Its Voronoi region is the intersection of open pseudodiscs and complements of pseudodiscs.
Claim 1. The intersection of n open pseudodiscs is simply connected[1] (or empty).
Claim 2. If each edge has the name of the pseudocircle which it is part of, then the edges along the boundary of the intersection of n pseudodiscs is a Davenport-Schinzel-Sequence of order 2.
Claim 3. The intersection of $n - 1$ open pseudodiscs and complements of pseudodiscs has $O(n)$ edges on its boundary. □

Theorem 3. *The Voronoi diagram based on a bisector system of pseudocircles fulfilling axioms (A2) and (A5) can be computed in expected time $O(n^2 \log n)$.*

6 Conclusion

Our algorithm covers many concrete Voronoi diagrams for which yet no efficient algorithm existed. The run time strongly depends on the size of m_j, the average number of components of a Voronoi region in a diagram of j sites. As we have seen, for pseudocircles this works very nicely and we obtain a nearly optimal run time, but there are examples were m_j is higher than the total complexity of the final diagram, thus in these cases our algorithm would be less efficient.

Also, if one wants to implement the algorithm one may not want to use Chazelles linear triangulation algorithm [9], but an easier approach, e. g. a randomized one [2]. This may be possible but it would require a new randomized analysis of the run time.

References

1. P.K. Agarwal, J. Pach and M. Sharir. State of the union, of geometric objects. Proc. Joint Summer Research Conf. on Discrete and Computational Geometry: 20 Years Later. Contemp. Math. 452, AMS, 2008, pp. 9–48.

[1] We call a set *simply connected* if it is both connected and has no holes.

2. N.M. Amato, M.T. Goodrich and E.A. Ramos. A Randomized Algorithm for Triangulating a Simple Polygon in Linear Time. Discrete & Computational Geometry. pp. 245–265, 2001.
3. F. Aurenhammer and H. Edelsbrunner. An Optimal Algorithm for Constructing the Weighted Voronoi Diagram in the Plane. Pattern Recognition 17(2), pp. 251–257, 1983.
4. F. Aurenhammer, R. Klein, and D.-T. Lee. Voronoi Diagrams and Delaunay Triangulations. World Scientific Publishing Company, 2013.
5. C. Bohler, P. Cheilaris, R. Klein, C.H. Liu, E. Papadopoulou, and M. Zavershynskyi. On the Complexity of Higher Order Abstract Voronoi Diagrams. Proceedings International Colloquium on Automata Languages and Programming (ICALP '13), pp. 208-219, 2013.
6. C. Bohler and R. Klein. Abstract Voronoi Diagrams with Disconnected Regions. Proceedings 24th International Symposium on Algorithms and Computation (ISAAC '13), pp. 306–316, 2013. To appear in International Journal of Computational Geometry & Applications.
7. C. Bohler, R. Klein, and C.H. Liu. Forest-Like Abstract Voronoi Diagrams in Linear Time. Proceedings 26th Canadian Conference on Computational Geometry (CCCG '14), 2014.
8. C. Bohler, C.H. Liu, E. Papadopoulou, and M. Zavershynskyi. A Randomized Divide and Conquer Algorithm for Higher-Order Abstract Voronoi Diagrams. Proceedings of the 25th International Symposium on Algorithms and Computations, pp. 27–37, Jeonju, Korea, 2014.
9. B. Chazelle. Triangulating a Simple Polygon in Linear Time. Discrete and Computational Geometry, 6, pp. 485–524, 1991.
10. O. Cheong, H. Everett, M. Glisse, J. Gudmundsson, S. Hornus, S. Lazard, M. Lee, and H-S. Na. Farthest-Polygon Voronoi Diagrams. Computational Geometry: Theory and Applications, 44 (4), pp. 234–247, 2011.
11. K. Clarkson, K. Mehlhorn, and R. Seidel. Four Results on Randomized Incremental Construction. Computational Geometry: Theory and Applications, 3, pp. 185–212, 1993.
12. K. Clarkson and P. Shor. Applications of Random Sampling in Computational Geometry, II. Discrete and Computational Geometry 4, pp. 387-421, 1989.
13. H. Edelsbrunner, L.J. Guibas, and J. Stolfi. Optimal Point Location in a Monotone Subdivision. SIAM Journal on Computing 15, pp. 317-340, 1986.
14. S. Har-Peled and B. Raichel. On the Complexity of Randomly Weighted Multiplicative Voronoi Diagrams. Proceedings 29th Annual Symposium on Computational Geometry (SoCG '13), 2013.
15. R. Klein, E. Langetepe, and Z. Nilforoushan. Abstract Voronoi Diagrams Revisited. Computational Geometry: Theory and Applications 42(9), pp. 885-902, 2009.
16. R. Klein and A. Lingas. Hamiltonian Abstract Voronoi Diagrams in Linear Time. Proceedings 5th International Symposium on Algorithms and Computation (ISAAC '94), pp. 11–19, 1994.
17. R. Klein, K. Mehlhorn, and St. Meiser. Randomized Incremental Construction of Abstract Voronoi Diagrams. Computational Geometry: Theory and Applications 3, pp. 157–184, 1993.
18. R. Seidel. A Simple and Fast Algorithm for Computing Trapezoidal Decompositions and for Triangulating Polygons. Computational Geometry: Theory and Applications 1, pp. 51–64, 1991.
19. M. Sharir and P. Agarwal. Davenport-Schinzel Sequences and Their Geometric Applications. Cambridge University Press, 1995.

7 Appendix

7.1 Proofs of Section 3

Proof. (of Lemma 3) We prove that if a trapezoid A of $\mathcal{H}(R)$ is intersected by t, then either a predecessor of A in $\mathcal{H}(R)$ is intersected by t, or A is detected in the search structure $\mathrm{SS}(F)$ of a face F, or found while tracing the region of t through a face F. In all these cases, we will detect A in the history graph.

So, assume that A is a trapezoid which was constructed during the insertion of a site $r_i \in R$, i. e. A is in $V^*(R_i)$ but not in $V^*(R_i \setminus \{r_i\})$. Now there are two cases.

Case 1: $A \subseteq \mathrm{VR}(r, R_i)$ and $r \neq r_i$.

Let $A_1, \ldots, A_k$ be the predecessors of A in $\mathcal{H}(R_i)$. Then $A \subseteq \bigcup_{i \in \{1,\ldots,k\}} \overline{A_i}$. Thus if A is intersected by t, then also at least one of its predecessors must be intersected by t.

Case 2: $A \subseteq \mathrm{VR}(r_i, R_i)$.

Let F be the face of $\mathrm{VR}(r_i, R_i)$ containing A.

If the region of t is not completely contained in F, then no face of the region of t may be an island in F. Thus there is a corridor contained in the region of t starting in A , crossing an edge e on the boundary of F and reaching into an adjacent trapezoid B. The trapezoid B must have been constructed during the insertion of r_i, thus the intersection of t and B is detected like in Case 1. Because the boundaries of B and F share a common edge e which is intersected by t, the region of t is traced through F^* and A is detected.

Now let the region of t be completely contained in F. Lets first observe that F is a face-vertex in $\mathcal{H}(R_i)$ and if $A_1, \ldots, A_k$ are its predecessors, then F is contained in $\bigcup_{i \in \{1,\ldots,k\}} \overline{A_i}$. Thus at least one of the predecessors of F is intersected by t and we reach F together with its search structure $\mathrm{SS}(F)$ while walking through $\mathcal{H}(R)$. Furthermore, because the region of t is contained in F the whole bisector $J(r_i, t)$ must be contained in F, thus for any point x on the bisector we will detect a trapezoid B contained in F, which is intersected by t. From there we start to trace the boundary of the region of t and reach A. □

7.2 Proofs of Section 4

Proof. (of Lemma 4) To test if a trapezoid A or edge e is intersected by t takes time $O(s)$, using our basic operations. Each trapezoid-vertex of $\mathcal{H}(R)$ has an outdegree of $O(s)$, because each trapezoid can be intersected in at most $O(s)$ components by a new Voronoi region. Thus for each intersected trapezoid-vertex at most $O(s)$ succeeding trapezoids are unsuccessfully tested for intersection with t. This takes time $O(s^2)$ per intersected trapezoid vertex.

Further, each trapezoid-vertex is linked to at most two edges e of the trapezoidal decomposition F^* for a face-vertex F. To test them for intersection with t or to detect that they already have a flag with respect to t takes time $O(s)$ per intersected trapezoid-vertex.

If such an edge e has been successfully tested for intersection with t, we have to test all trapezoids bordering on e from inside of F for intersection with t. There may be many such trapezoids, but fortunately, if e is intersected by t, then all trapezoids bordering on e are in conflict with t and thus they can be upper bounded by c. Within the face F, each trapezoid has at most 4 neighbors, thus while continuing the tracing within the face F, for each intersected trapezoid at most four adjacent trapezoids may be intersected unsuccessfully for intersection with t. The flags prevent us from testing edges and trapezoids several times for intersection with t. Thus altogether this process takes time $O(sc)$.

For each face-vertex F of $\mathcal{H}(R)$ reached, we have to perform a point location searching in $\mathrm{SS}(F)$ which takes time $O(\log(F))$. A very rough (but sufficient for O-notations) estimation of $|F|$ is $O(sn^2)$, which gives us $O(\log(F)) \in O(\log(\max\{s,n\}))$. If the searching was successful, we trace the region of t and again for each intersected trapezoid we test at most 4 additional trapezoids unsuccessfully.

Each trapezoid intersected by t is also in conflict with t, thus the overall running time to compute E_t is in $O(s^2 c \log(\max\{s,n\}))$. □

Proof. (of Lemma 5) Let A be a trapezoid-vertex of the history graph $\mathcal{H}(r_1,\ldots,r_j)$, where the sites $r_1,\ldots,r_j$ are inserted in this order. Now let $t = r_{j+1} \in S \setminus R$ randomly chosen. If A is in conflict with t, then A with its description $D_R(A)$ would never have been constructed, if t would have been inserted as the first site, i. e., $A \notin \mathcal{H}(t, r_1,\ldots,r_j)$. By elementary set calculus, it follows:

$$\begin{aligned}&|\mathcal{H}(r_1,\ldots,r_j) \setminus \mathcal{H}(t,r_1,\ldots,r_j)|\\ =&|\mathcal{H}(r_1,\ldots,r_j)| - |\mathcal{H}(t,r_1,\ldots,r_j)| + |\mathcal{H}(t,r_1,\ldots,r_j) \setminus \mathcal{H}(r_1,\ldots,r_j)|.\end{aligned}$$

The expectation of $|\mathcal{H}(r_1,\ldots,r_j)| - |\mathcal{H}(t,r_1,\ldots,r_j)|$ is ≤ 0, because $|\{r_1,\ldots,r_j\}| < |\{t,r_1,\ldots,r_j\}|$ implies for randomized inputs the expectation $|\mathcal{H}(r_1,\ldots,r_j)| \le |\mathcal{H}(t,r_1,\ldots,r_j)|$. Thus the equation above has an expected value

$$\le E(\underbrace{|\mathcal{H}(t,r_1,\ldots,r_j) \setminus \mathcal{H}(r_1,\ldots,r_j)|}_{=:X}).$$

If $A \in |\mathcal{H}(t,r_1,\ldots,r_j) \setminus \mathcal{H}(r_1,\ldots,r_j)|$, then $t \in \mathrm{set}(D_{R\cup\{t\}}(A))$. Let i be minimal with $A \in V^*(t,r_1,\ldots,r_i)$ and thus $r_i \in \mathrm{set}(D_{R\cup\{t\}}(A))$.
Then the size of X is upper bounded by

$$\le \sum_{i=2}^{j} \underbrace{|\{A \in V^*(t,r_1,\ldots,r_i) | t, r_i \in \mathrm{set}(D_{R_i\cup\{t\}}(A))\}|}_{=:Y},$$

where $R_i = \{r_1, \ldots, r_i\}$. By choosing t and r_i randomly from $\{t, r_1, \ldots, r_i\}$ we can estimate the expectation of Y by

$$\leq \frac{1}{(i+1)i} \sum_{(x,y)\in\{t,r_1,\ldots,r_i\}^2, x\neq y} |\{A \in V^*(t, r_1, \ldots, r_i) | x, y \in \text{set}(D_{R_i\cup\{t\}}(A))\}|$$

and since $\text{set}(D_R(A))$ consists of at most 16 different sites it follows

$$\leq \frac{1}{(i+1)i} \cdot 16 \cdot 15 \cdot \underbrace{|V^*(t, r_1, \ldots, r_i)|}_{\leq C|\text{V}(t,r_1,\ldots,r_i)| \leq C m_{i+1}(i+1)} \leq C\frac{m_{i+1}}{i}$$

for a constant C. Thus the expectation of X is

$$\leq \sum_{i=2}^{j} C\frac{m_{i+1}}{i} \in O\left(\sum_{i=2}^{j} \frac{m_i}{i}\right).$$

This implies $c \in O(\sum_{i=2}^{j} \frac{m_i}{i})$. □

Proof. (of Lemma 6) Let $R_i = \{r_1, \ldots, r_i\}$ be the set of the first i sites inserted in this order. The size of $\mathcal{H}(S)$ equals $\sum_{i=3}^{n} \mathcal{H}(R_i) \setminus \mathcal{H}(R_{i-1})$. If $x \in \mathcal{H}(R_i) \setminus \mathcal{H}(R_{i-1})$, then by the definition of the history graph, either x is a a face of $\text{VR}(r_i, R_i)$ or x is a trapezoid of $V^*(R_i) \setminus V^*(R_{i-1})$. Because the number of faces in the history graph is less than the number of trapezoids (each face-vertex has at least one unique succeeding trapezoid), the number of all vertices of the history graph is $\leq$ two times the number of trapezoid-vertices. If x is a trapezoid-vertex, then $r_i \in D_{R_i}(x)$ which implies:

$$E(\mathcal{H}(R_i) \setminus \mathcal{H}(R_{i-1})) \leq \frac{2}{i} \sum_{r\in R_i} \sum_{A\in\text{trapezoids}(V^*(R_i))} \delta(A, r),$$

$$\text{where } \delta(A, r) = \begin{cases} 1, & \text{if } r \in \text{set}(D_{R_i}(A)) \\ 0, & \text{else.} \end{cases}$$

$$\leq \frac{2}{i}(16|V^*(R_i)|) \in O(m_i).$$

Summing up over all i from $i = 3$ to n, we get $E(|\mathcal{H}(S)|) \in O(\sum_{i=3}^{n} m_i)$. □

Proof. (of Lemma 7) The set E_t equals the set of all trapezoids of $V^*(R_j) \setminus V^*(R_j \cup \{t\})$. With the same arguments as in the two previous proofs one can show that the expected size of this set is $O(m_j)$. □

7.3 Comments about Section 5

Lemma 8. *If $|S| = 3$, then all Voronoi regions are connected.*

Proof. Let $S = \{p, q, r\}$. Then the region of p equals the intersection of the two dominance regions $D(p, q) \cap D(p, r)$. A dominance region is either an open pseudodisc or the open complement of a pseudodisc. It is easy to see that the intersection of two dominance regions is always connected. □

For unbounded bisecting curves we have the property that if for all subsets S' of S of size 3 the Voronoi regions in $V(S')$ are connected, then all Voronoi regions in $V(S)$ are connected, see [15, Lemma 14]. This is something which is no longer true, if bisectors may be closed curves, as has been shown already for multiplicatively weighted Voronoi Diagrams [3]. Here one region may consist of $O(n)$ faces and the total complexity of $V(S)$ is $O(n^2)$. Interestingly there is no connection at all between the property that for all subsets S' of S of size k, $1 \leq k \leq n-1$, all regions in $V(S')$ are connected and the property that all regions in $V(S)$ are connected.

Lemma 9. *Let S be a set of n sites. We can choose a set of bisecting curves such that for all proper subsets $S' \subsetneq S$ all Voronoi regions in $V(S')$ are connected but there is a region in $V(S)$ which is disconnected.*

Proof. An example for such a system of bisecting curves is depicted in Figure 7. Let $S = \{p_1, \ldots, p_n\}$. The uppermost circle is the bisector $J(p_1, p_2), J(p_2, p_3), \ldots, J(p_2, p_n)$, where $D(p_1, p_2), D(p_3, p_2) \ldots, D(p_n, p_2)$ equal the outer face of the circle. The circle to the right of the uppermost one is the bisector $J(p_1, p_3)$, $J(p_3, p_4), \ldots, J(p_3, p_n)$, where $D(p_1, p_3), D(p_4, p_3) \ldots, D(p_n, p_3)$ equal the outer face of it. Finally the circle to the left of the uppermost one is the bisector $J(p_1, p_n)$, where $D(p_1, p_n)$ is the outer face.

Now the region of p_1 equals the plane minus the union of the circles. The uppermost disc is the region of p_2, the disc to the right minus its left one is the region of p_3 and so on. Thus in $V(S)$, where all circles are present, the region of p_1 is disconnected into two faces. But as soon as one circle is missing, i. e. when we have a subset S' not containing one of the sites $p_2, \ldots, p_n$, then the region of p_1 grows together through this circle and becomes connected in $V(S')$. All other regions are connected in all $V(S')$, where S' is a subset of S. □

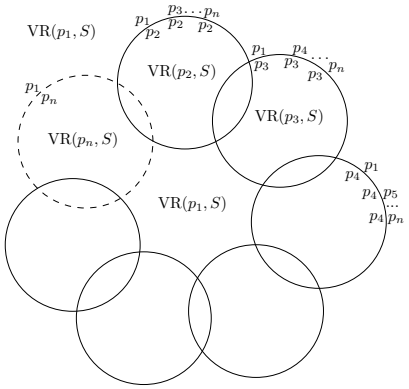

Fig. 7: An example where the region of p_1 is disconnected in $V(S)$, but for proper subsets $S' \subsetneq S$ all regions in $V(S')$ are connected.

Proof. (of Theorem 2)

Let p be a site in S. Its Voronoi region is the intersection of open pseudodiscs and complements of pseudodiscs.

Claim 1. The intersection of n open pseudodiscs is simply connected[2] (or empty).

Proof. It is clear that the intersection has no holes, otherwise one of the pseudodiscs would have a hole. Thus, it remains to show that it is connected, which follows from [15, Lemma 14], because if all Voronoi regions in any diagram of three sites are connected, then related unbounded bisectors may intersect in at most two points, and these intersections are transversal. Thus, the bisectors can be connected to closed ones at infinity giving us a system of pseudocircles. Now the claim follows from the lemma. □

Because of Claim 1, the boundary of the intersection of n pseudodiscs is a closed curve consisting of pseudocircle segments. We call these segments edges.

Claim 2. If each edge has the name of the pseudocircle which it is part of, then the edges along the boundary of the intersection of n pseudodiscs is a Davenport-Schinzel-Sequence of order 2.

Proof. Let an edge have the name q if it belongs to the pseudocircle $J(p, q)$. Now suppose there are edges q and r alternating three times in the sequence, i. e. $q \dots r \dots q \dots r$. Then the two bisectors $J(p, q)$ and $J(p, r)$ must intersect in more than two points, a contradiction. □

This claim shows that there can be at most $2n - 1$ edges on the boundary of the intersection of n pseudodiscs.

Now recall that $\mathrm{VR}(p, S)$ is the intersection of both open pseudodiscs and complements of pseudodiscs. Let F be the intersection of the pseudodiscs $D(p, q_i)$. By Claim 2, F is connected and has $O(n)$ edges on its boundary. It is clear that $\mathrm{VR}(p, S)$ is contained in F. Place a vertex on each edge on the boundary of F and for each pseudodisc $D(q, p)$ (whose complement is a dominance region of p) intersecting F, place a vertex q in $D(q, p) \cap F$; observe that each pseudodisc can intersect F in at most one connected component.

Connect two vertices q and r belonging to $D(q, p)$ and $D(r, p)$ by an edge iff $\partial D(q, p)$ and $\partial D(r, p)$ intersect in a point in F not contained in any other $D(s, p)$, i. e. $J(q, p)$ and $J(r, p)$ intersect in a vertex on the boundary of the region of p Further, connected a vertex q belonging to $D(q, p)$ and a vertex v belonging to an edge e on the boundary of F iff $\partial D(q, p)$ intersects e in a point not contained in any other $D(r, p)$. Finally, connect two vertices belonging to edges e and e' on the boundary of F, if e and e' are incident.

Now consider the resulting graph G, see Figure 8. Because two pseudocircles can intersect in at most two points, G is a planar graph without parallel edges.

[2] We call a set *simply connected* if it is both connected and has no holes.

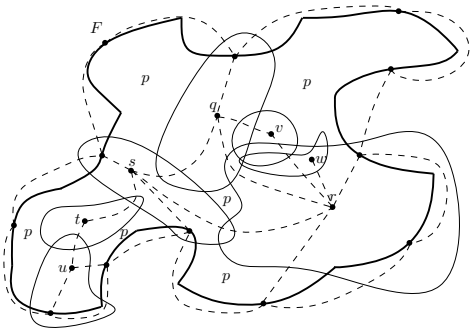

Fig. 8: Proof of the last part of Theorem 2. Fat curves depict the boundary of F, normal curves the boundaries of the pseudodiscs $D(q,p)$ and dashed curves edges of the graph G.

Further we can draw the edges e between two vertices $q \in D(q,p)$ and $r \in D(r,p)$ such that $e \subseteq D(q,p) \cup D(r,p)$, which is possible, because by definition e exists only if $D(q,p) \cap D(r,p) \neq \emptyset$. Then it is clear that each connected component of $\mathrm{VR}(p,S)$ corresponds to a unique face of G (but G may have more faces than $\mathrm{VR}(p,S)$). Because G has $O(n)$ vertices it also has $O(n)$ faces and edges, which proves the upper bound of the Theorem.

An example, where the bounds are tight is the Voronoi diagram of multiplicatively weighted points, see [3]. □